arXiv:1511.06845v1 [math.NA] 21 Nov 2015

OPINS: An Orthogonally Projected Implicit Null-Space Method for Singular and Nonsingular Saddle-Point Systems*

Cao Lu† Tristan Delaney† Xiangmin Jiao†‡

November 24, 2015

Abstract

Saddle-point systems appear in many scientific and engineering applications. The systems can be sparse, symmetric or nonsymmetric, and possibly singular. In many of these applications, the number of constraints is relatively small compared to the number of unknowns. The traditional null-space method is inefficient for these systems, because it is expensive to find the null space explicitly. Some alternatives, notably constraint-preconditioned or projected Krylov methods, are relatively efficient, but they can suffer from numerical instability and even nonconvergence. In addition, most existing methods require the system to be nonsingular or be reducible to a nonsingular system. In this paper, we propose a new method, called *OPINS*, for singular and nonsingular saddle-point systems. OPINS is equivalent to the null-space method with an orthogonal projector, without forming the orthogonal basis of the null space explicitly. OPINS can not only solve for the unique solution for nonsingular saddle-point problems, but also find the minimum-norm solution in terms of the solution variables for singular systems. The method is efficient and easy to implement using existing Krylov solvers for singular systems. At the same time, it is more stable than the other alternatives, such as projected Krylov methods. We present some preconditioners to accelerate the convergence of OPINS for nonsingular systems, and compare OPINS against some present state-of-the-art methods for various types of singular and nonsingular systems.

[1]This work was supported by DoD-ARO under contract #W911NF0910306. The third author is also supported by a subcontract to Stony Brook University from Argonne National Laboratory under Contract DE-AC02-06CH11357 for the SciDAC program funded by the Office of Science, Advanced Scientific Computing Research of the U.S. Department of Energy.

[2]Dept. of Applied Math. & Stat., Stony Brook University, Stony Brook, NY 11794, USA.

[3]Corresponding author. Email: jiao@ams.sunysb.edu.

1 Introduction

We consider the numerical solution of the following saddle-point system

$$\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B}^T \\ \boldsymbol{B} & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{y} \end{bmatrix} = \begin{bmatrix} \boldsymbol{f} \\ \boldsymbol{g} \end{bmatrix}, \tag{1}$$

where $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ may be nonsymmetric, $\boldsymbol{B} \in \mathbb{R}^{m\times n}$, $\boldsymbol{x}, \boldsymbol{f} \in \mathbb{R}^n$, and $\boldsymbol{y}, \boldsymbol{g} \in \mathbb{R}^m$. The coefficient matrix, denoted by $\boldsymbol{K}$, is in general nonsymmetric and indefinite. It may be nonsingular or singular, and often very ill-conditioned even when it is nonsingular.

This type of saddle-point problems arises in many scientific applications. For example, it can arise from solving the following constrained optimization problem,

$$\min_{\boldsymbol{x}} \frac{1}{2}\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} - \boldsymbol{f}^T\boldsymbol{x} \quad \text{subject to} \quad \boldsymbol{B}\boldsymbol{x} = \boldsymbol{g}, \tag{2}$$

using the method of Lagrange multipliers, where $\boldsymbol{A}$ is the Hessian of the quadratic objective function, and $\boldsymbol{B}\boldsymbol{x} = \boldsymbol{g}$ defines a constraint hyperplane for the minimization. In this case, $\boldsymbol{A}$ is typically symmetric, $\boldsymbol{x}$ contains the *optimization variables* or *solution variables*, and $\boldsymbol{y}$ contains the *Lagrange multipliers* or *constraint variables*. The optimality conditions are referred as Karush-Kuhn-Tucker conditions, and the system (1) is often called the *KKT system* [24], with $\boldsymbol{K}$ being the *KKT matrix*. In PDE discretization, the Lagrange multipliers are often used to enforce nodal conditions, such as sliding boundary conditions and continuity constraints at hanging nodes [15].

Solving the saddle-point problems is particularly challenging. One may attempt to solve it using a general-purpose solver, such as a direct solver or an iterative solver [9, 13, 31], or the range-space method [32]. For large and sparse saddle-point systems, a more powerful method is the null-space method [6], which solves for $\boldsymbol{x}$ first and then $\boldsymbol{y}$, with an iterative method as its core solver. The null-space method computes $\boldsymbol{x}$ by first finding a particular solution, and then solving its complementary component in null($\boldsymbol{B}$) (c.f. Algorithm 1). This method is physically and geometrically meaningful, as it effectively finds a solution $\boldsymbol{x}_*$ within the constraint hyperplane such that $\boldsymbol{A}\boldsymbol{x}_* - \boldsymbol{f} \in \text{range}(\boldsymbol{B}^T)$. This is satisfied if and only if $\boldsymbol{f} \in \text{range}(\boldsymbol{A}) + \text{range}(\boldsymbol{B}^T)$. Let $\boldsymbol{Z}$ denote the matrix composed of a set of basis vectors of null ($\boldsymbol{B}$), and we say $\boldsymbol{Z}$ is *orthonormal* if the basis is orthonormal. The traditional null-space method is efficient if null($\boldsymbol{B}$) is low dimensional, since it involves a small linear system with the coefficient matrix $\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}$. In exact arithmetic, $\boldsymbol{Z}$ does not need to be orthonormal. However, with rounding errors, it is desirable for $\boldsymbol{Z}$ to be orthonormal (or nearly orthonormal), so that the condition number of $\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}$ is close to that of $\boldsymbol{A}$. However, if the dimension of null($\boldsymbol{B}$) is high, i.e., when $m \ll n$, determining such a nearly orthonormal basis is computationally expensive, making the traditional null-space method impractical for such applications.

Another class of methods, which we refer to as *implicit null-space methods*, are effective for large-scale saddle-point systems. Examples of this method include the Krylov subspace methods with a constraint preconditioner [30, 18] and

the projected Krylov methods (KSP) [19]. In exact arithmetic, these methods are equivalent to the null-space method with an orthonormal basis of null($\boldsymbol{B}$), but they do not require computing the basis explicitly. These methods are usually more efficient than directly applying the Krylov subspace methods to the whole system. However, with rounding errors, these methods may suffer from numerical instability, due to loss of orthogonality. The instability may be mitigated by iterative refinements [18, 19], but there is no guarantee that the solver would not stagnate even with iterative refinements.

In this paper, we propose an implicit orthogonal projection method, called *OPINS*, which is a more stable variant of implicit null-space methods. Specifically, we compute an orthonormal basis for range $\left(\boldsymbol{B}^T\right)$, which can be constructed efficiently when $m \ll n$ using a stable algorithm such as QR factorization with column pivoting (QRCP). Let $\boldsymbol{U}$ denote the matrix composed of such an orthonormal basis. Instead of using $\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}$ in the null-space method, we use the orthogonal projector $\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp} \equiv \boldsymbol{I} - \boldsymbol{U}\boldsymbol{U}^T$ to construct a singular but compatible linear system, with a symmetric coefficient matrix $\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{A}\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}$. We solve this system using a solver for singular systems, such as MINRES [25, 11] and SYMMLQ [25] for symmetric systems, and GMRES [29, 27], LSQR [26] or LSMR [14] for nonsymmetric systems. We also propose preconditioners for saddle-point systems, based on the work in [19]. The resulting OPINS method is highly efficient when $m \ll n$, and it is stable, robust, and easy to implement using existing Krylov-subspace method for singular systems, without the issue of stagnation suffered by other implicit null-space methods.

Besides being more stable, another advantage of OPINS is that it can be applied to singular saddle-point systems, which arise in various applications. For example, in the application of fracture mechanics, a solid object may contain many cracks and isolated pieces, and an elasticity model is in general underconstrained, leading to singular saddle-point systems. In the literature, most methods assume the saddle-point system is nonsingular. When this assumption is violated, for example when there are redundant constraints or insufficient constraints, these methods may force the user to add artificial boundary conditions or soft constraints to make the system nonsingular, which unfortunately may undermine the physical accuracy of the solution. We consider singular saddle-point systems under the assumption that $\boldsymbol{f} \in \text{range}(\boldsymbol{A}) + \text{range}(\boldsymbol{B}^T)$, in other words, the saddle-point system is compatible in terms of $\boldsymbol{f}$. For problems arising from constrained minimization, such a system finds the solution that minimizes $\|\boldsymbol{x}\|$ among all the solutions of

$$\min_{\boldsymbol{x}} \frac{1}{2}\boldsymbol{x}^T\boldsymbol{A}\boldsymbol{x} - \boldsymbol{f}^T\boldsymbol{x} \quad \text{subject to} \quad \min_{\boldsymbol{x}} \|\boldsymbol{g} - \boldsymbol{B}\boldsymbol{x}\|. \tag{3}$$

Throughout this paper, the norms are in 2-norm unless otherwise noted. We refer to the above solution as the *minimum-norm solution* of a singular saddle-point system. Geometrically, this effectively defines a constraint hyperplane in a least-squares sense, and then require the objective function to be minimized exactly within the hyperplane, and it defines a unique solution under the assump-

tion of $\boldsymbol{f} \in \text{range}(\boldsymbol{A}) + \text{range}(\boldsymbol{B}^T)$. Conversely, if $\boldsymbol{f} \notin \text{range}(\boldsymbol{A}) + \text{range}(\boldsymbol{B}^T)$, the objective function does not have a critical point in the constraint hyperplane, so the constrained minimization is ill-posed. OPINS delivers a more accurate and systematic method to find the minimum-norm solution for such singular saddle-point problems.

The remainder of the paper is organized as follows. Section 2 reviews some background and related work. Section 3 introduces the Orthogonally Projected Implicit Null-Space Method (OPINS) for singular and nonsingular saddle-point systems. Section 4 introduces effective preconditioners for OPINS. Section 5 presents some numerical results with OPINS as well as comparisons against some state-of-the-art methods. Section 6 concludes the paper with a discussion on future research directions.

2 Background

We briefly review some important concepts and related methods for solving saddle-point problems. There is a large body of literature on saddle-point problems; see [6] for a comprehensive survey up to early 2000s and the references in [19] for more recent works. We focus our discussions on the null-space methods, since they are the most relevant to our proposed method. For completeness, we will also briefly discuss some other methods for saddle-point systems and singular linear systems.

2.1 Explicit Null-Space Methods

One of the most powerful methods for solving saddle-point systems is the null-space method. Geometrically, for a nonsingular saddle-point problem arising from constrained minimization, this method finds a critical point of an objective function within a constraint hyperplane $\boldsymbol{B}\boldsymbol{x} = \boldsymbol{g}$. Algorithm 1 outlines the algorithm, where $\boldsymbol{x}_p \in \mathbb{R}^n$ denotes a particular solution within the constraint hyperplane. Let q denote the rank of the constraint matrix $\boldsymbol{B}$, and $q \leq m$. The column vectors of the $\boldsymbol{Z} \in \mathbb{R}^{n\times(n-q)}$ form a basis of null($\boldsymbol{B}$), and $\boldsymbol{B}\boldsymbol{Z} = \boldsymbol{0}$. Step 3 finds a component $\boldsymbol{x}_n$ in null($\boldsymbol{B}$), i.e. a vector tangent to the constraint hyperplane, so that $\boldsymbol{x}_n + \boldsymbol{x}_p$ is at a critical point. After determining $\boldsymbol{x}$, the final step finds the Lagrange multipliers $\boldsymbol{y}$ in range($\boldsymbol{B}$).

In the algorithm, step 3 is the most critical, which solves the equation

$$\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{v} = \boldsymbol{Z}^T\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right), \tag{4}$$

within null($\boldsymbol{B}$). We refer to (4) as the *null-space equation*, and denote its coefficient matrix by $\hat{\boldsymbol{N}}$. This matrix is $(n-q)\times(n-q)$, which is smaller than the original matrix $\boldsymbol{K}$. When $\boldsymbol{A}$ is symmetric and positive semidefinite and $\text{null}\,(\boldsymbol{A}) \cap \text{null}\,(\boldsymbol{B}) = \{\boldsymbol{0}\}$, $\hat{\boldsymbol{N}}$ is SPD [6], and (4) can be solved efficiently using preconditioned conjugate gradient (CG). However, if $\hat{\boldsymbol{N}}$ is symmetric but indefinite or is nonsymmetric, then an alternative iterative solver, such MINRES,

Algorithm 1 Null-Space Method

input: $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{f}$, $\boldsymbol{g}$, tolerance for the iterative solver (if used)
output: $\boldsymbol{x}$, $\boldsymbol{y}$ (optional)

1: solve $\boldsymbol{B}\boldsymbol{x}_p = \boldsymbol{g}$
2: compute $\boldsymbol{Z}$, composed of basis vectors of null($\boldsymbol{B}$)
3: solve $\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{v} = \boldsymbol{Z}^T(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p)$
4: $\boldsymbol{x}_n \leftarrow \boldsymbol{Z}\boldsymbol{v}$ and $\boldsymbol{x} \leftarrow \boldsymbol{x}_p + \boldsymbol{x}_n$
5: $\boldsymbol{y} \leftarrow \boldsymbol{B}^{+T}(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x})$

SYMMLQ [25], or GMRES [29], can be used; see textbooks such as [28] for details of these iterative solvers.

In the traditional null-space method, the matrix $\boldsymbol{Z}$ is constructed explicitly. We refer to such an approach as the *explicit null-space method*. In exact arithmetic, it is not necessary for $\boldsymbol{Z}$ to be orthonormal. If $\boldsymbol{B}$ has full rank, then there exists an $n \times n$ permutation matrix $\boldsymbol{P}$ such as $\boldsymbol{B}\boldsymbol{P} = [\boldsymbol{B}_1 \mid \boldsymbol{B}_2]$, where $\boldsymbol{B}_1$ is an $m \times m$ nonsingular matrix, and $\boldsymbol{B}_2$ is $m \times (n-m)$. Then, one could simply choose $\boldsymbol{Z}$ to be [16]

$$\boldsymbol{Z} = \boldsymbol{P}\begin{bmatrix} -\boldsymbol{B}_1^{-1}\boldsymbol{B}_2 \\ \boldsymbol{I} \end{bmatrix},$$

where $\boldsymbol{I}$ is the $(n-m) \times (n-m)$ identity matrix. However, if the columns of $\boldsymbol{Z}$ are too far from being orthonormal, $\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}$ may be ill-conditioned, which in turn can cause slow convergence for iterative solvers or large errors in the resulting solution. Therefore, it is desirable for $\boldsymbol{Z}$ to be orthonormal or nearly orthonormal, so that $\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}$ (approximately) preserves the condition number of $\boldsymbol{A}$. If the dimension of null($\boldsymbol{B}$) is high, i.e., when $m \ll n$, $\boldsymbol{Z}$ is typically quite dense. Determining an orthonormal $\boldsymbol{Z}$ requires the full QR factorization of $\boldsymbol{B}^T$, which takes $\mathcal{O}(n^3)$ operations. Therefore, explicit null-space methods are impractical for large-scale applications.

2.2 Implicit Null-Space Methods

For saddle-point systems where $m \ll n$, it is desirable to avoid constructing an orthonormal basis of null($\boldsymbol{B}$) explicitly. One such approach is to use a Krylov subspace method with a *constraint preconditioner* [18, 22], which has the form

$$\boldsymbol{M} = \begin{bmatrix} \boldsymbol{G} & \boldsymbol{B}^T \\ \boldsymbol{B} & \boldsymbol{0} \end{bmatrix}, \tag{5}$$

where $\boldsymbol{G}$ is an approximation of $\boldsymbol{A}$. This preconditioner is indefinite, and it was shown in [30] that it provides optimal bounds for the maximum eigenvalues among similar indefinite preconditioners. Since the preconditioner is indefinite, it is not obvious whether we can use it as a preconditioner for methods such as CG or MINRES, which typically require symmetric positive-definite preconditioners. As shown in [18, 19], one can apply the preconditioner to the modified

saddle-point system

$$\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B}^T \\ \boldsymbol{B} & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_n \\ \boldsymbol{y} \end{bmatrix} = \begin{bmatrix} \boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p \\ \boldsymbol{0} \end{bmatrix}, \tag{6}$$

with preconditioned CG and MINRES. The vector $\boldsymbol{x}_n$ in (6) is the same as that from step 4 in Algorithm 1, but the vector $\boldsymbol{y}$ computed from (6) may be inaccurate, so one may need to solve for $\boldsymbol{y}$ separately once $\boldsymbol{x}$ is obtained. In exact arithmetic, the Krylov subspace method with this constraint preconditioner is equivalent to the *projected Krylov method* in [19]. Because of this equivalence, we will use the names *projected Krylov methods* and *Krylov subspace methods with constraint preconditioning* interchangeably in this paper, although these methods somewhat differ in their implementation details.

The projected Krylov method is closely related to the preconditioned null-space method on the null-space equation (4) with an orthonormal $\boldsymbol{Z}$. The stability of the method requires the computed $\boldsymbol{x}_n$ to be exactly in null($\boldsymbol{B}$) at each step. However, the rounding errors can quickly introduce a nonnegligible component in range($\boldsymbol{B}^T$), which can cause the method to break down [18, 19]. To mitigate this issue, the constraint preconditioner must be applied "exactly," by solving the preconditioned system with a direct method followed by one or more steps of iterative refinement per iteration [18, 19]. The iterative refinement introduces extra cost, and there is no guarantee that it would recover orthogonality between $\boldsymbol{x}_n$ and range($\boldsymbol{B}$) to machine precision, so the projected Krylov method may still stagnate. Besides its potential instability, the projected Krylov methods typically assume the KKT system is nonsingular [19]. It is desirable to develop a more stable version of the implicit null-space method that can also be applied to singular KKT systems. The OPINS method proposed in this paper achieves this goal.

2.3 Other Methods for Saddle-Point Systems

Besides the null-space methods, another class of methods for saddle-point problems is the range-space method [32]. It first obtains $\boldsymbol{y}$ by solving the system

$$\left(\boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{B}^T\right)\boldsymbol{y} = \boldsymbol{B}\boldsymbol{A}^{-1}\boldsymbol{f} - \boldsymbol{g}, \tag{7}$$

where the coefficient matrix is the Schur complement, and then computes $\boldsymbol{x}$ by solving

$$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{f} - \boldsymbol{B}^T\boldsymbol{y}. \tag{8}$$

The range-space method can be attractive if a factorization of $\boldsymbol{A}$ is available. However, computing the Schur complement is expensive if $\boldsymbol{A}$ is large and sparse, and the method is not directly applicable if $\boldsymbol{A}$ is singular.

The null-space and range-space methods both leverage the special structures of the saddle-point systems. In some cases, one may attempt to solve the whole system (1) directly using a factorization-based method, such as LDL^T decomposition for symmetric systems [31]. These methods are prohibitively

expensive for large-scale problems. In addition, since the solution variables $\boldsymbol{x}$ and constraint variables $\boldsymbol{y}$ have different physical meanings, the entries in $\boldsymbol{A}$ and $\boldsymbol{B}$ may have very different scales. As a result, $\boldsymbol{K}$ may be arbitrarily ill-conditioned, and these methods may break down or produce inaccurate solutions due to poor scaling. Another class of method is iterative solvers with preconditioners [7]. The constraint preconditioner is a special case, which is equivalent to an implicit null-space method, with some optimality properties among similar indefinite preconditioners [3, 30]. Some other preconditioners include the block diagonal [13], block triangular [9], and multigrid [1]. These methods do not distinguish between $\boldsymbol{x}$ and $\boldsymbol{y}$, so they tend to have more difficulties when $\boldsymbol{K}$ is singular or ill-conditioned.

2.4 Methods for Singular Saddle-Point Systems

The methods we discussed above typically assume nonsingular saddle-point systems. For general singular saddle-point problems, one may resort to truncated SVD [17] or rank-revealing QR, which are computationally expensive. One may also apply an iterative solver for singular systems, such as MINRES and SYMMLQ [25] for compatible symmetric systems, MINRES-QLP [11] for incompatible symmetric systems, Breakdown-free GMRES [27], LSQR [26] and LSMR [14] for nonsymmetric systems, and GMRES [29] for compatible nonsymmetric systems with $\text{range}(\boldsymbol{K}) = \text{range}(\boldsymbol{K}^T)$. Without preconditioners, except for Breakdown-free GMRES [27], these methods can find the minimum-norm solution of compatible singular systems when they are applicable. However, these methods do not distinguish between $\boldsymbol{x}$ and $\boldsymbol{y}$ in the solver. As a result, these methods minimize $\|\boldsymbol{x}\|^2 + \|\boldsymbol{y}\|^2$, which may substantially differ from the minimum-norm solution of $\boldsymbol{x}$ as defined in (3), especially when $\boldsymbol{K}$ is ill-conditioned. OPINS will overcome this issue by leveraging the iterative solvers in the implicit null-space method in a stable and efficient fashion, and in turn offer a new effective method for solving singular saddle-point systems.

3 Proposed Method

In this section, we introduce the *Orthogonally Projected Implicit Null-Space Method*, or *OPINS*, for solving saddle-point systems (1). Similar to the other implicit null-space methods, OPINS does not require the explicit construction of the basis of $\text{null}(\boldsymbol{B})$, and it is particularly effective when $m \ll n$. However, unlike previous null-space methods, OPINS enforces orthogonality explicitly and hence enjoys better stability. It is also applicable to singular saddle-point systems. In the following subsections, we present the OPINS method, its derivation, and the analysis of its cost and stability.

Algorithm 2 OPINS: Orthogonally Projected Implicit Null-Space Method

input: $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{f}$, $\boldsymbol{g}$, tolerances for rank estimation and iterative solver
output: $\boldsymbol{x}$, $\boldsymbol{y}$ (optional)

1: $\boldsymbol{B}^T\boldsymbol{P} = \boldsymbol{Q}\boldsymbol{R}$ {QR factorization with column pivoting}
2: $\boldsymbol{U} \leftarrow \boldsymbol{Q}_{1:q}$, where q is rank($\boldsymbol{B}$) estimated from QRCP
3: $\boldsymbol{x}_p \leftarrow \boldsymbol{B}^{+}\boldsymbol{g} = \boldsymbol{U}\boldsymbol{R}_{1:q,1:q}^{-T}\left(\boldsymbol{P}^T\boldsymbol{g}\right)_{1:q}$
4: solve $\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{A}\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{w} = \boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right)$ using iterative singular solver
5: $\boldsymbol{x} \leftarrow \boldsymbol{x}_p + \boldsymbol{x}_n$, where $\boldsymbol{x}_n = \boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{w}$
6: $\boldsymbol{y} \leftarrow \boldsymbol{B}^{+T}\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}\right) = \boldsymbol{P}_{:,1:q}\boldsymbol{R}_{1:q,1:q}^{-1}\boldsymbol{U}^T\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}\right)$

3.1 Algorithm Description

A core idea of OPINS is to use the orthogonal projector onto null($\boldsymbol{B}$), constructed from an orthonormal basis of range($\boldsymbol{B}^T$). Let $\boldsymbol{U}$ denote the matrix composed of an orthonormal basis of range($\boldsymbol{B}^T$), which can be computed using truncated SVD or QR with column pivoting (QRCP) [10, 17]. Since $\boldsymbol{U}$ is orthonormal, $\boldsymbol{\Pi}_{\boldsymbol{U}} \equiv \boldsymbol{U}\boldsymbol{U}^T$ is the unique orthogonal projector onto range($\boldsymbol{B}^T$), and

$$\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp} \equiv \boldsymbol{I} - \boldsymbol{\Pi}_{\boldsymbol{U}} = \boldsymbol{I} - \boldsymbol{U}\boldsymbol{U}^T$$

is its complementary orthogonal projector onto null($\boldsymbol{B}$). Note that $\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp} = \boldsymbol{\Pi}_{\boldsymbol{Z}}$, where $\boldsymbol{Z}$ is composed of an orthonormal basis of null($\boldsymbol{B}$). Let $q = \text{rank}(\boldsymbol{B})$. Using this projector, OPINS replaces the $(n-q)\times(n-q)$ null-space equation (4) in the null space method with the $n \times n$ singular system

$$\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{A}\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{w} = \boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right), \tag{9}$$

and then solves it using a solver for singular systems. We refer to the above equation as the *projected null-space (PNS) equation*, and denote its coefficient matrix as $\boldsymbol{N}$.

Algorithm 2 outlines the complete OPINS algorithm, which applies to non-singular or compatible singular systems. The first two steps find an orthonormal basis of $\boldsymbol{B}^T$ using QRCP, where $\boldsymbol{P}$ is an $m \times m$ permutation matrix, so that the diagonal values of $\boldsymbol{R}$ are sorted in descending order. $\boldsymbol{Q} \in \mathbb{R}^{n\times m}$ is orthonormal, and $\boldsymbol{R} \in \mathbb{R}^{m\times m}$ is upper triangular. For stability, QRCP should be computed based on Householder QR factorization [17]. If $\boldsymbol{B}$ is rank deficient, its rank can be estimated from the magnitude of the diagonal entries in $\boldsymbol{R}$, or more robustly using a condition-number estimator [17]. The first q columns of $\boldsymbol{Q}$ form an orthonormal basis of range $\left(\boldsymbol{B}^T\right)$. With QRCP, both $\boldsymbol{x}_p$ in step 3 and $\boldsymbol{y}$ in step 6 can also be solved efficiently. When $\boldsymbol{B}$ is rank deficient, so is $\boldsymbol{R}$. We use $\boldsymbol{R}_{1:q,1:q}^{-T}$ and $\boldsymbol{R}_{1:q,1:q}^{-1}$ to denote the forward and back substitutions on $\boldsymbol{R}_{1:q,1:q}$.

The key step of the algorithm is the solution of the PNS equation (9) in step 4, which is singular. As we will show in Section 3.2, (9) is compatible under the assumption that (1) is compatible in terms of $\boldsymbol{f}$, i.e., $\boldsymbol{f} \in \text{range}(\boldsymbol{A}) +$

range($\boldsymbol{B}^T$). Therefore, we can solve it using a Krylov subspace method for compatible singular systems, as we discuss in more detail in Subsection 3.2.3. A key operation in these methods is the multiplication of the coefficient matrix with a vector. For the multiplication with $\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}$ with any vector $\boldsymbol{v} \in \mathbb{R}^n$,

$$\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{v} = \boldsymbol{v} - \boldsymbol{U}\left(\boldsymbol{U}^T\boldsymbol{v}\right),$$

which can be computed stably and efficiently. Note that $\boldsymbol{U}$ is not stored explicitly either, but as a collection of Householder reflection vectors in QRCP.

3.2 Detailed Derivation of OPINS

The derivation of OPINS is similar to that of the null-space method. For completeness, we will start with the derivation of the explicit null-space method, and then extend it to derive OPINS. We will also discuss the solution techniques of the PNS equation.

3.2.1 Null-Space Method for Singular and Nonsingular Systems

The null-space method is typically derived algebraically for nonsingular saddle-point systems. Since we also consider singular systems that may be partially incompatible, it is instructive to consider an alternative but equivalent derivation, which has clear geometric meanings for singular saddle-point problems in the context of constraint minimization (3).

Suppose (1) is a saddle-point system compatible in terms of $\boldsymbol{f}$, i.e., $\boldsymbol{f} \in$ range($\boldsymbol{A}$) + range($\boldsymbol{B}^T$). Then, there exists $\boldsymbol{x}_* \in \mathbb{R}^n$ and $\boldsymbol{y}_* \in \mathbb{R}^m$, such that

$$\boldsymbol{A}\boldsymbol{x}_* - \boldsymbol{f} = -\boldsymbol{B}^T\boldsymbol{y}_* \text{ subject to } \min_{\boldsymbol{x}_*} \|\boldsymbol{g} - \boldsymbol{B}\boldsymbol{x}_*\|. \tag{10}$$

The constraint defines a hyperplane, in the least squares sense, whose tangent space is orthogonal to $\boldsymbol{B}^T$. The vector $\boldsymbol{A}\boldsymbol{x}_* - \boldsymbol{f}$ corresponds to the gradient of a nonlinear objective function. Equation (10) indicates that $\boldsymbol{A}\boldsymbol{x}_* - \boldsymbol{f}$ is a linear combination of the column vectors of $\boldsymbol{B}^T$ with the Lagrange multipliers $\boldsymbol{y}_*$ as the coefficients. In other words, $\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_*$ is orthogonal to the constraint hyperplane. Hence, we can rewrite (10) as $\boldsymbol{Z}^T(\boldsymbol{A}\boldsymbol{x}_* - \boldsymbol{f}) = \boldsymbol{0}$, where $\boldsymbol{Z}$ is a basis of null($\boldsymbol{B}$). The saddle-point system is then equivalent to

$$\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{x}_* = \boldsymbol{Z}^T\boldsymbol{f} \text{ subject to } \min_{\boldsymbol{x}_*} \|\boldsymbol{g} - \boldsymbol{B}\boldsymbol{x}_*\|,$$

which is an alternative statement of the constraint minimization.

Let $\boldsymbol{x}_* = \boldsymbol{x}_n + \boldsymbol{x}_p$, where $\boldsymbol{x}_p \in$ range($\boldsymbol{B}^T$) is a point on the constraint hyperplane. There exists $\boldsymbol{x}_n = \boldsymbol{Z}\boldsymbol{v} \in$ null($\boldsymbol{B}$), i.e. a vector tangent to the constraint hyperplane, such that

$$\boldsymbol{Z}^T\boldsymbol{A}(\boldsymbol{x}_n + \boldsymbol{x}_p) = \boldsymbol{Z}^T\boldsymbol{f},$$

or equivalently,

$$\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{v} = \boldsymbol{Z}^T\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_{\boldsymbol{p}}\right).$$

This is the null-space equation (4). If $\boldsymbol{f} \in \text{range}(\boldsymbol{A}) + \text{range}(\boldsymbol{B}^T)$, then the null-space equation (4) is compatible.

In the above derivation, $\boldsymbol{Z}$ does not need to be orthonormal. However, if $\boldsymbol{Z}$ is far from being orthonormal, the system (4) may have a large condition number, which may affect the convergence of the iterative solvers and the accuracy of the numerical solution.

3.2.2 OPINS for Singular and Nonsingular Systems

To derive OPINS, now assume $\boldsymbol{Z}$ is orthonormal. We further multiply $\boldsymbol{Z}$ to both sides of the null-space equation and then obtain

$$\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{v} = \boldsymbol{Z}\boldsymbol{Z}^T\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right). \tag{11}$$

In addition, we rewrite $\boldsymbol{x}_n = \boldsymbol{Z}\boldsymbol{v}$ as the orthogonal projection of a vector $\boldsymbol{w} \in \mathbb{R}^n$ onto null $(\boldsymbol{B})$, i.e.,

$$\boldsymbol{Z}\boldsymbol{v} = \boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{w}.$$

Substituting it into (11), we have

$$\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{w} = \boldsymbol{Z}\boldsymbol{Z}^T\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right), \tag{12}$$

which is equivalent to (9) in step 4 of Algorithm 2.

The above transformation may seem counterintuitive, as we intentionally constructed a singular system (9), which is larger than the null-space equation (4). However, (9) has three key properties: First, it uses an orthogonal projector to ensure that $\boldsymbol{x}_n$ is exactly in null$(\boldsymbol{B})$, and hence it overcomes the instability associated with the projected Krylov methods. Second, since $\boldsymbol{\Pi}_{\boldsymbol{Z}} = \boldsymbol{Z}\boldsymbol{Z}^T = \boldsymbol{I} - \boldsymbol{U}\boldsymbol{U}^T = \boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}$, we can compute the projection by finding $\boldsymbol{U}$, which is much more efficient than finding $\boldsymbol{Z}$ when $q \leq m \ll n$. Third, since (9) is always singular, and QRCP in step 1 supports rank-deficient $\boldsymbol{B}$, we can apply OPINS to a saddle-point system regardless of whether $\boldsymbol{A}$, $\boldsymbol{B}$, or $\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}$ is singular.

Although the PNS system is singular, it is in fact compatible, due to the following property.

Lemma 1. *If the saddle-point system (1) is compatible in terms of $\boldsymbol{f}$, i.e., $\boldsymbol{f} \in$ range$(\boldsymbol{A})$ + range$(\boldsymbol{B}^T)$, then the null-space equation (4) and the PNS equation (9) are both compatible.*

The proof of this lemma directly follows from the derivations above. Furthermore, if the null-space equation (4) is nonsingular, we have the following lemma.

Lemma 2. *If the null-space equation (4) is nonsingular, then the coefficient matrix of the PNS system (9) has the same row and column spaces, i.e., range$(\boldsymbol{N}) =$ range$(\boldsymbol{N}^T)$.*

The proof is straightforward since $\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}$ has the same row and column spaces. With the compatibility assumption, Lemma 2 ensures that the standard

GMRES can be applied to system (9) without breaking down [27, 29]. In the following subsection, we will address the efficient solution of these compatible systems that may be singular or nonsingular.

3.2.3 Solution of Orthogonally Projected Null-Space Equation

Since the PNS system (4) has an infinite number of solutions, a natural question is which solution of the system suffices in producing the minimum-norm solution in the sense of (3). In the following, we will address the question first for systems with a nonsingular null-space equation, followed by more general singular systems.

Theorem 3. *Given a saddle-point system (1) where the null-space equation (4) is nonsingular and the solution to (1) is unique in terms of $\boldsymbol{x}$, then the solution of (9) recovers this unique $\boldsymbol{x}$.*

Proof. Consider the alternative form of the PNS equation in (12). The system is compatible, so we can always find a solution $\boldsymbol{w} \in \mathbb{R}^n$ for the equality to hold. Since $\boldsymbol{Z}^T\boldsymbol{Z} = \boldsymbol{I}$, left-multiplying $\boldsymbol{Z}^T$ on both sides of (12), we obtain the PNS equation

$$\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{w} = \boldsymbol{Z}^T\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right).$$

Since system (4) is nonsingular, $\boldsymbol{Z}^T\boldsymbol{w}$ recovers the unique solution for $\boldsymbol{v}$, and in turn recovers the unique $\boldsymbol{x}_n$ and $\boldsymbol{x}$. □

Note that nonsingular (4) includes the cases where the saddle-point system (1) is nonsingular. However, it does not necessarily imply that (1) is nonsingular, because $\boldsymbol{y}$ may not be unique. An implication of Theorem 3 is that we have the flexibility of solving (4) with any solver even if (1) is singular.

For a general singular saddle-point system, the situation is more complicated. Assume (1) is compatible in terms of $\boldsymbol{f}$, the following theorem indicates that OPINS finds the minimum-norm solution of $\boldsymbol{x}$.

Theorem 4. *Given a saddle-point system (1) compatible in terms of $\boldsymbol{f}$, i.e., $\boldsymbol{f} \in$ range$(\boldsymbol{A})$ + range$(\boldsymbol{B}^T)$; if $\boldsymbol{x}$ is the minimum-norm solution of the PNS equation (9), then $\|\boldsymbol{x}\|$ is minimized among all the solutions $\boldsymbol{x} = \boldsymbol{x}_n + \boldsymbol{x}_p$ that satisfy the constraint* $\min_{\boldsymbol{x}} \|\boldsymbol{g} - \boldsymbol{B}\boldsymbol{x}\|$.

Proof. Note that $\boldsymbol{x} = \boldsymbol{x}_p + \boldsymbol{x}_n$, where $\boldsymbol{x}_p \in \text{range}(\boldsymbol{B}^T)$ and $\boldsymbol{x}_n \in \text{null}(\boldsymbol{B})$, so $\|\boldsymbol{x}\|^2 = \|\boldsymbol{x}_p\|^2 + \|\boldsymbol{x}_n\|^2$. In step 3 of QRCP, $\boldsymbol{x}_p$ is the minimum-norm solution in range$(\boldsymbol{B}^T)$ that minimizes $\|\boldsymbol{g}-\boldsymbol{B}\boldsymbol{x}\|$. Therefore, we only need to show that $\|\boldsymbol{x}_n\|$ is minimized in null$(\boldsymbol{B})$, where $\boldsymbol{x}_n = \boldsymbol{Z}\boldsymbol{v}$. This is satisfied if $\|\boldsymbol{v}\|$ is minimized among the exact solutions to the null-space equation (4), i.e.,

$$\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{v} = \boldsymbol{Z}^T(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p).$$

Since $\boldsymbol{Z}$ is orthonormal, $\|\boldsymbol{x}_n\| = \|\boldsymbol{v}\|$. In OPINS, if $\boldsymbol{w}$ is an exact solution to (9),

$$\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{w} = \boldsymbol{Z}\boldsymbol{Z}^T(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p).$$

Since $\boldsymbol{Z}^T\boldsymbol{Z} = \boldsymbol{I}$, by left-multiplying $\boldsymbol{Z}^T$ on both sides, we have

$$\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{w} = \boldsymbol{Z}^T(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p),$$

so $\boldsymbol{v} = \boldsymbol{Z}^T\boldsymbol{w}$ is an exact solution of the null-space equation (4). Note that $\|\boldsymbol{v}\| = \|\boldsymbol{Z}^T\boldsymbol{w}\| \leq \|\boldsymbol{Z}^T\|\|\boldsymbol{w}\| = \|\boldsymbol{w}\|$, and it is an equality if $\boldsymbol{w} \in \text{range}(\boldsymbol{Z}) = \text{null}(\boldsymbol{B})$, i.e., $\boldsymbol{w} = \boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{w} = \boldsymbol{x}_n$. Therefore, the minimum-norm solution of $\boldsymbol{w}$ in (9) minimizes $\|\boldsymbol{x}_n\|$. Furthermore, $\boldsymbol{g} - \boldsymbol{B}\boldsymbol{x} = \boldsymbol{g} - \boldsymbol{B}\boldsymbol{x}_p$, so $\|\boldsymbol{x}\|$ is minimized among all solutions under the constraint $\min_{\boldsymbol{x}} \|\boldsymbol{g} - \boldsymbol{B}\boldsymbol{x}\|$. □

Note that Theorem 4 is more general than Theorem 3, as it includes Theorem 4 as a special case. An implication of the two theorems is that if the null-space equation is nonsingular, we can use any singular solver for OPINS. Examples of such solvers include MINRES [25] and SYMMLQ [25] for symmetric and compatible systems, GMRES [29], LSQR [26] or LSMR [14] for nonsymmetric systems. In addition, we can use any left, right, or symmetric preconditioner for the solver. However, if the null-space equation is singular, we must use a solver that can compute the minimum-norm solution for compatible singular systems, which may exclude GMRES for nonsymmetric systems if $\text{range}(\boldsymbol{A}) \neq \text{range}(\boldsymbol{A}^T)$. In addition, we can use only left preconditioners that do not alter the null space of the coefficient matrix. We will further discuss the preconditioners in Section 3.2.3.

Finally, regarding the $\boldsymbol{y}$ component in (4), the values of $\boldsymbol{y}$ are immaterial for many applications. However, if desired, we can obtain $\boldsymbol{y}$ by solving $\boldsymbol{B}^T\boldsymbol{y} = (\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x})$ using QRCP. In general, $\|\boldsymbol{y}\|$ may not be minimized if $\boldsymbol{B}$ is rank deficient, but the norm is typically small.

3.3 Efficiency of OPINS

We now analyze the computational cost of OPINS. There are two components that are relatively more expensive. The first is the QRCP in step 1, which takes $\mathcal{O}(\frac{4}{3}m^2n)$ operations when using Householder transformation. When $m \ll n$, this operation is far more efficient than finding an orthonormal basis of null($\boldsymbol{B}$). After obtaining the QR factorization, steps 3, 5, and 6 all take $\mathcal{O}(mn)$ operations. In addition, if $\boldsymbol{B}$ is sparse, $\boldsymbol{Q}$ and $\boldsymbol{R}$ are in general also sparse, leading to even more cost savings. The other one is the solution of the singular system (9). When $\boldsymbol{A}$ is large and sparse, it is not advisable to use truncated SVD or rank-revealing QR factorization for this system. Instead, we apply a Krylov-subspace method for singular systems. Within each iteration of these methods, the dominating operation is matrix-vector multiplications, which cost $\mathcal{O}(N + nq)$, where N denote the total number of nonzeros in $\boldsymbol{A}$. The convergence of these methods depend on the nonzero eigenvalues [20]. The following proposition correlates the eigenvalues of the coefficient matrices $\hat{\boldsymbol{N}}$ and $\boldsymbol{N}$ in (4) and (9), respectively.

Proposition 5. *Given a saddle-point system (1), the PNS matrix in (9) has the same nonzero eigenvalues as the null-space system (4) with an orthonormal* $\boldsymbol{Z}$.

Proof. Let $\hat{\boldsymbol{N}} = \boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}$ and $\boldsymbol{N} = \boldsymbol{Z}\left(\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\right)\boldsymbol{Z}^T = \boldsymbol{Z}\hat{\boldsymbol{N}}\boldsymbol{Z}^T$. If $\hat{\lambda}$ is a nonzero eigenvalue of $\hat{\boldsymbol{N}}$, and $\hat{\boldsymbol{x}}$ is a corresponding eigenvalue, then

$$\boldsymbol{N}(\boldsymbol{Z}\hat{\boldsymbol{x}}) = \left(\boldsymbol{Z}\hat{\boldsymbol{N}}\boldsymbol{Z}^T\right)(\boldsymbol{Z}\hat{\boldsymbol{x}}) = \boldsymbol{Z}\hat{\boldsymbol{N}}\hat{\boldsymbol{x}} = \hat{\lambda}\boldsymbol{Z}\hat{\boldsymbol{x}},$$

so $\hat{\lambda}$ and $\boldsymbol{Z}\hat{\boldsymbol{x}}$ form a pair of eigenvalue and eigenvector of $\boldsymbol{N}$. Conversely, if λ a nonzero eigenvalue of $\boldsymbol{N}$ and $\boldsymbol{x}$ is its corresponding eigenvalue, then

$$\hat{\boldsymbol{N}}(\boldsymbol{Z}^T\boldsymbol{x}) = \left(\boldsymbol{Z}^T\boldsymbol{Z}\right)\left(\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\right)\boldsymbol{Z}^T\boldsymbol{x} = \boldsymbol{Z}^T\boldsymbol{N}\boldsymbol{x} = \lambda\boldsymbol{Z}^T\boldsymbol{x}.$$

Therefore, $\boldsymbol{N}$ and $\hat{\boldsymbol{N}}$ have the same nonzero eigenvalues. □

Based on the above proposition, the convergence rate of the Krylov subspace method on the PNS equation (9) is identical to that on the null-space equation (4) with an orthonormal $\boldsymbol{Z}$. However, to accelerate the convergence of these methods, it is desirable to use preconditioners, which we discuss next.

4 Preconditioners for OPINS

In OPINS, the most time consuming step is typically the iterative solver for the PNS equation (9) in step 4. To speed up its computation, it is critical to use preconditioners. In this section, we present some principles for constructing effective preconditioners for PNS equations, based on the recent work on the projected Krylov methods [19].

Let $\boldsymbol{N} = \boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{A}\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}$. The general idea of preconditioning is to find a matrix $\boldsymbol{M}$ that approximates the coefficient matrix $\boldsymbol{N}$, or $\boldsymbol{M}^+$ that approximates the pseudoinverse $\boldsymbol{N}^+$. The latter form is more convenient for solving the PNS systems. Algorithm 3 outlines the pseudocode of the preconditioned OPINS with a left preconditioner. The preconditioning routine takes an operator $\boldsymbol{M}$ to evaluate $\boldsymbol{M}^+\boldsymbol{b}$ for any $\boldsymbol{b} \in \mathbb{R}^n$. Note that for symmetric systems, most preconditioned Krylov-subspace methods would apply the preconditioners symmetrically. Specifically, suppose $\boldsymbol{M}^+$ has a symmetric factorization $\boldsymbol{M}^+ = \boldsymbol{L}\boldsymbol{L}^T$, then these methods solve the equation

$$\boldsymbol{L}^T\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{A}\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{L}\tilde{\boldsymbol{w}} = \boldsymbol{L}^T\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right)$$

in the preconditioned method, and then computes $\boldsymbol{w} = \boldsymbol{L}\tilde{\boldsymbol{w}}$. Typically, the algorithm is constructed such that the explicit factorization $\boldsymbol{M}^+ = \boldsymbol{L}\boldsymbol{L}^T$ is not needed. We omit the details of such preconditioned Krylov-subspace methods; interested readers may refer to [4, 28].

In this section, we will focus on the construction of $\boldsymbol{M}$ or $\boldsymbol{M}^+$. A straightforward choice of $\boldsymbol{M}$ is the approximation of $\boldsymbol{A}$. Possible candidates include SSOR-type preconditioners, incomplete factorization, and multigrid methods. We propose to approximate $\boldsymbol{N}^+$ with $\boldsymbol{M}^+ = \boldsymbol{P}_G = \boldsymbol{Z}\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}\boldsymbol{Z}^T$, where

Algorithm 3 Preconditioned OPINS

input: $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{f}$, $\boldsymbol{g}$, $\boldsymbol{G}$, tolerances for rank estimation and iterative solver
output: $\boldsymbol{x}$, $\boldsymbol{y}$ (optional)

1: do first three steps of Algorithm 2
2: solve $\boldsymbol{M}^{+}\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{A}\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{w} = \boldsymbol{M}^{+}\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p)$ using a preconditioned Krylov-subspace method
3: do the last two steps of Algorithm 2.

$\boldsymbol{G}$ is an approximation of $\boldsymbol{A}$ and $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ is nonsingular. We refer to this preconditioner as the *projected preconditioner*. As we shall show in this section, for nonsingular systems it is equivalent to the constraint preconditioner in the projected Krylov methods [19, 22]. In the following, we will first discuss the details of implementing $\boldsymbol{P}_G$ as an operator for nonsymmetric systems, and then specialize it for symmetric systems.

4.1 Preconditioners for Nonsymmetric Systems

If the saddle-point system is nonsymmetric, the projected preconditioner applies as long as $\boldsymbol{G}$ is nonsingular. For example, we can take $\boldsymbol{G}$ as the SOR-style preconditioner or incomplete LU factorization of $\boldsymbol{A}$. Then, $\boldsymbol{P}_G = \boldsymbol{Z}\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}\boldsymbol{Z}^T$, and we use $\boldsymbol{P}_G$ as a left preconditioner. The following proposition states that OPINS with the projected preconditioner is equivalent to applying $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ as a left preconditioner to a null-space equation, and hence is equivalent to the constraint-preconditioned null-space methods [19, 22].

Proposition 6. *Assume $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ is nonsingular. OPINS with the projected preconditioner as a left preconditioner is equivalent to left preconditioning the null-space equation with $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$.*

Proof. Since $\boldsymbol{P}_G$ is applied as a left preconditioner, the preconditioned OPINS solves

$$\underbrace{\boldsymbol{Z}\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}\boldsymbol{Z}^T}_{\boldsymbol{P}_G}\underbrace{\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{Z}^T}_{\boldsymbol{N}}\boldsymbol{w} = \underbrace{\boldsymbol{Z}\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}\boldsymbol{Z}^T}_{\boldsymbol{P}_G}\boldsymbol{Z}\boldsymbol{Z}^T(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p).$$

Since $\boldsymbol{Z}^T\boldsymbol{Z} = \boldsymbol{I}$ and $\boldsymbol{Z}$ is orthonormal, it is equivalent to solving

$$\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{v} = \left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}\boldsymbol{Z}^T(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p),$$

i.e., applying $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ as a left preconditioner to (4), where $\boldsymbol{v} = \boldsymbol{Z}^T\boldsymbol{w}$. □

As shown in [22], the eigenvalues of the constraint-preconditioned null-space equation are well clustered if $\boldsymbol{G}$ is a good approximation of $\boldsymbol{A}$. Due to the above equivalence, the projected preconditioner is a good choice for OPINS. Note that when applying $\boldsymbol{P}_G$ as a left preconditioner, it does not alter the null-space of $\boldsymbol{N}$, as claimed in the following proposition.

Algorithm 4 Operator $\boldsymbol{P}_G$ for Projected Preconditioner

input: $\boldsymbol{G}$, $\boldsymbol{U}$, $\boldsymbol{b}$
output: $\boldsymbol{s} = \boldsymbol{P}_G\boldsymbol{b}$

1: solve $\boldsymbol{G}\boldsymbol{r} = \boldsymbol{b}$
2: solve $(\boldsymbol{U}^T\boldsymbol{G}^{-1}\boldsymbol{U})\boldsymbol{t} = \boldsymbol{U}^T\boldsymbol{r}$
3: solve $\boldsymbol{G}\boldsymbol{s} = \boldsymbol{b} - \boldsymbol{U}\boldsymbol{t}$

Proposition 7. *Assume $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ is nonsingular. OPINS with the projected preconditioner as the left preconditioner does not alter the null-space of $\boldsymbol{N}$, i.e., null$(\boldsymbol{P}_G\boldsymbol{N})$ = null$(\boldsymbol{N})$.*

Proof. For any vector $\boldsymbol{v} \in \text{null}(\boldsymbol{N})$, $\boldsymbol{v} \in \text{null}(\boldsymbol{P}_G\boldsymbol{N})$. If $\boldsymbol{v} \in \text{null}(\boldsymbol{P}_G\boldsymbol{N})$, we have

$$\boldsymbol{Z}\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{v} = 0.$$

Since $\boldsymbol{Z}$ has full column rank and $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ is nonsingular, $\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{v} = 0$. It follows that $\boldsymbol{v} \in \text{null}(\boldsymbol{N})$. □

The above proposition indicates that we can apply the projected preconditioner as a left preconditioner to singular saddle-point that is compatible in $\boldsymbol{f}$, while ensuring $\boldsymbol{x}$ has minimum norm.

The remaining task is to find a way to provide $\boldsymbol{P}_G$ as an operator for efficient computation of $\boldsymbol{s} = \boldsymbol{P}_G\boldsymbol{b}$ for any $\boldsymbol{b} \in \mathbb{R}^n$. Note that $\boldsymbol{s}$ is the solution to the following modified but simpler saddle-point system

$$\begin{bmatrix} \boldsymbol{G} & \boldsymbol{U} \\ \boldsymbol{U}^T & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{s} \\ \boldsymbol{t} \end{bmatrix} = \begin{bmatrix} \boldsymbol{b} \\ \boldsymbol{0} \end{bmatrix}, \tag{13}$$

where $\boldsymbol{U}$ is composed of an orthonormal basis of range$(\boldsymbol{B}^T)$. Because from the null-space method, we have

$$\boldsymbol{s} = \boldsymbol{Z}\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}\boldsymbol{Z}^T\boldsymbol{b} = \boldsymbol{P}_G\boldsymbol{b}. \tag{14}$$

This implies that $\boldsymbol{P}_G$ can be given as an operator through the solution of (13). (5) is similar to the procedure for evaluating the constraint preconditioner in the projected Krylov method [19], for which the off-diagonal entries are $\boldsymbol{B}^T$ and $\boldsymbol{B}$ instead of $\boldsymbol{U}$ and $\boldsymbol{U}^T$. By replacing $\boldsymbol{B}^T$ by $\boldsymbol{U}$, (5) can be solved more efficiently, and it is also applicable if $\boldsymbol{B}$ is rank deficient.

If $\boldsymbol{G}$ is a simple matrix or operator, such as the preconditioner based on SOR or incomplete factorization, we can solve (5) efficiently by using the range-space method, given by (7) and (8), especially when $m \ll n$. Algorithm 4 outlines the procedure for computing $\boldsymbol{P}_G\boldsymbol{b}$ using the range-space method, where $\boldsymbol{U}$ is composed of an orthonormal basis of $\boldsymbol{B}^T$.

4.2 Preconditioners for Symmetric Systems

For symmetric systems, most preconditioned Krylov subspace methods require the preconditioner to be symmetric. Hence, we require that $\boldsymbol{G}$ is symmetric and $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ is SPD. Therefore, $\boldsymbol{P}_G$ is symmetric and positive semi-definite. If the preconditioner is applied as a left preconditioner, the result for nonsymmetric systems also apply to symmetric systems. Assuming the system is symmetric, the following proposition states that OPINS with the projected preconditioner is equivalent to applying $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ as a preconditioner in solving the corresponding null-space equation, regardless whether it is applied as a left preconditioner or a symmetric preconditioner.

Proposition 8. *Assume $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ is SPD and the saddle-point system is symmetric. OPINS with the projected preconditioner is equivalent to applying $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ as the preconditioner for solving the null-space equation in the null-space method.*

Proof. If $\boldsymbol{P}_G$ is applied as a left preconditioner, the proof is the same as for nonsymmetric systems. Now consider $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ is SPD, then so is $\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}$, which has a Cholesky factorization

$$\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1} = \boldsymbol{L}_G\boldsymbol{L}_G^T.$$

Then $\boldsymbol{P}_G = \boldsymbol{Z}\left(\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}\right)^{-1}\boldsymbol{Z}^T$ has a symmetric factorization

$$\boldsymbol{P}_G = \boldsymbol{L}\boldsymbol{L}^T = \boldsymbol{Z}\boldsymbol{L}_G(\boldsymbol{Z}\boldsymbol{L}_G)^T = \boldsymbol{Z}\boldsymbol{L}_G\boldsymbol{L}_G^T\boldsymbol{Z}^T,$$

where $\boldsymbol{L} = \boldsymbol{Z}\boldsymbol{L}_G$. If $\boldsymbol{P}_G$ is applied symmetrically, the preconditioned OPINS would solve the equation

$$\underbrace{\boldsymbol{L}_G^T\boldsymbol{Z}^T}_{\boldsymbol{L}^T}\underbrace{\boldsymbol{Z}\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}\boldsymbol{Z}^T}_{\boldsymbol{N}}\underbrace{\boldsymbol{Z}\boldsymbol{L}_G}_{\boldsymbol{L}}\tilde{\boldsymbol{w}} = \underbrace{\boldsymbol{L}_G^T\boldsymbol{Z}^T}_{\boldsymbol{L}^T}\boldsymbol{Z}\boldsymbol{Z}^T\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right),$$

where $\boldsymbol{Z}^T\boldsymbol{Z} = \boldsymbol{I}$ and $\boldsymbol{w} = \boldsymbol{L}_G\tilde{\boldsymbol{w}}$. Therefore, it is equivalent to solving

$$\boldsymbol{L}_G^T\underbrace{\boldsymbol{Z}^T\boldsymbol{A}\boldsymbol{Z}}_{\hat{\boldsymbol{N}}}\boldsymbol{L}_G\tilde{\boldsymbol{w}} = \boldsymbol{L}_G^T\boldsymbol{Z}^T\left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right),$$

which is equivalent to applying $\boldsymbol{Z}^T\boldsymbol{G}\boldsymbol{Z}$ as a symmetric preconditioner to (4). □

For preconditioned Krylov subspace methods for symmetric systems, $\boldsymbol{P}_G$ can be computed as an operator using Algorithm 4. Note that if the system is singular, applying the preconditioner symmetrically may alter the null space of the coefficient matrix, unless it is equivalent to applying a left preconditioner. Therefore, additional care must be taken to find the minimum-norm solution when applying a preconditioner symmetrically.

Table 1: Summary of test problems.

problem	len($\boldsymbol{x}$)	len($\boldsymbol{y}$)	rank($\boldsymbol{K}$)	source
3d-var	3240	3	3243	Analysis for climate modeling [21]
sherman5	3312	20	3332	nonsymmetric problem from [8]
mosarqp1	2500	700	3200	quadratic programming [23]
fracture	780	92	786	2-D elasticity with fracture [5]
can_61	61	20	81	symmetric problem from [8]
random	100	20	120	random nonsingular matrix
random-s	100	20	90	random singular matrix

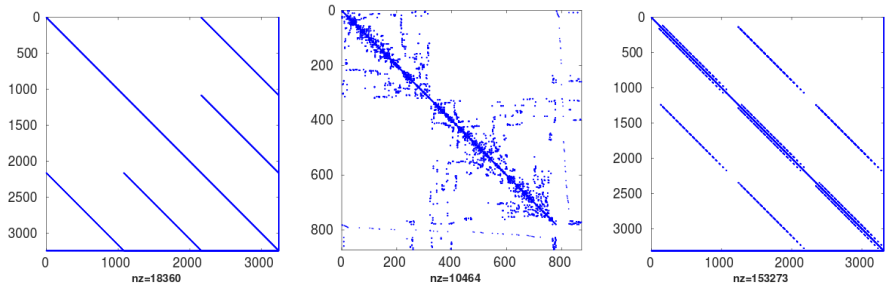

(a) 3d-var (b) fracture (c) sherman5

Figure 1: Sparsity patterns of KKT matrices in 3d-var, fracture and sherman5.

5 Numerical Results

In this section, we evaluate OPINS with various test problems, including both singular and nonsingular systems. The experiments are mainly focused on symmetric systems. Some results of nonsymmetric systems are also included in Section 5.1. We start by evaluating the performance of OPINS with and without preconditioners, to demonstrate the importance of preconditioners and the effectiveness of the projected preconditioner. We then compare OPINS against some present state-of-the-art methods for symmetric nonsingular and singular systems. We use a few test problems, as summarized in Table 1 arising from constrained minimization, finite element analysis, climate modeling, and random matrices. Among these problems, the first six are sparse, and Figure 1 shows the sparsity patterns for some of their KKT matrices. If the right-hand sides were unavailable, we generate them by multiplying the matrix with a random vector. For all problems, we set the convergence tolerance to 10^{-10} for the residual in iterative methods, and set the tolerance for QRCP to 10^{-12}.

In terms of error measures, different methods use different convergence criteria internally. For a direct comparison, we present the convergence results in two difference measures. The first is the residual of $\boldsymbol{x}$ within the null space of $\boldsymbol{B}$, i.e.,

$$\boldsymbol{r} := \boldsymbol{\Pi}_{\boldsymbol{Z}} (\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p) - \boldsymbol{\Pi}_{\boldsymbol{Z}} \boldsymbol{A}\boldsymbol{x}_n,$$

where $\boldsymbol{x}_p$ is the particular solution in range($\boldsymbol{B}^T$) and $\boldsymbol{x}_n$ is the corresponding solution in null($\boldsymbol{B}$). In the context of constraint minimization, this residual is an indicator of how well $\nabla_{\boldsymbol{d}}\phi = \boldsymbol{0}$ is satisfied for the objective function ϕ for all directions within the constraint hyperplane. To make the metric scale independent, we measure the residual relative to the right-hand side of (4), i.e.,

$$\text{relative residual in } \boldsymbol{x} := \|\boldsymbol{r}\| \,/\, \left\|\boldsymbol{\Pi}_{\boldsymbol{Z}} \left(\boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}_p\right)\right\|.$$

When comparing OPINS with methods that solve for $\boldsymbol{x}$ and $\boldsymbol{y}$ simultaneously, such as preconditioned Krylov methods, we calculate the residual in $\boldsymbol{x}$ by computing $\boldsymbol{U}$ using QR factorization. For a more complete comparison, in addition to the above error metric, we also compute the residual for the whole system (1) relative to the right-hand side in terms of both $\boldsymbol{x}$ and $\boldsymbol{y}$.

5.1 Effectiveness of Preconditioners

For Krylov subspace methods, the preconditioners have significant impact on the convergence rate. First let us consider symmetric systems. The core solver in OPINS is based on MINRES, so we assess the effectiveness of OPINS with a straightforward preconditioner as well as the projected preconditioner as described in Section 4. For simplicity, we choose $\boldsymbol{G}$ as the Jacobi preconditioner for both preconditioners, and denote them as OPINS-J and OPINS-P, respectively.

We solve the test case 3d-var with unpreconditioned OPINS, OPINS-J and OPINS-P. Figure 2(a) shows the convergence results measured in terms of the $\boldsymbol{x}$ residual. For 3d-var, the block $\boldsymbol{A}$ is nearly diagonal and is strongly diagonal dominant, so both preconditioners worked well and performed significantly better than unpreconditioned OPINS. Between the preconditioners, the projected preconditioner performed better than the Jacobi preconditioner for MINRES alone, because the projected preconditioner provides a better approximation to the whole matrix. This indicates that the projected preconditioner is effective for accelerating OPINS, if $\boldsymbol{G}$ is a good approximation of $\boldsymbol{A}$.

To demonstrate the applicability to nonsymmetric systems, we solve the problem sherman5 from the matrix market database [8]. $\boldsymbol{A}$ comes from an oil reservoir simulation and $\boldsymbol{B}$ is a random matrix. Since the system is nonsymmetric, we choose $\boldsymbol{G}$ as the ILU factorization of $\boldsymbol{A}$ for both preconditioners, and denote the two strategies as OPINS-ILU and OPINS-P, respectively. The inner solver is GMRES with the number of restarts set to 50. Figure 2(b) shows the convergence results for the three approaches. The results show that OPINS indeed works for nonsymmetric systems. Similar to the symmetric case, OPINS-P is faster than OPINS-ILU while both are much faster than unpreconditioned OPINS.

5.2 Assessment for Symmetric Nonsingular Systems

We now perform a more in-depth assessment of OPINS for symmetric nonsingular systems. Since unpreconditioned OPINS is not effective for nonsingular

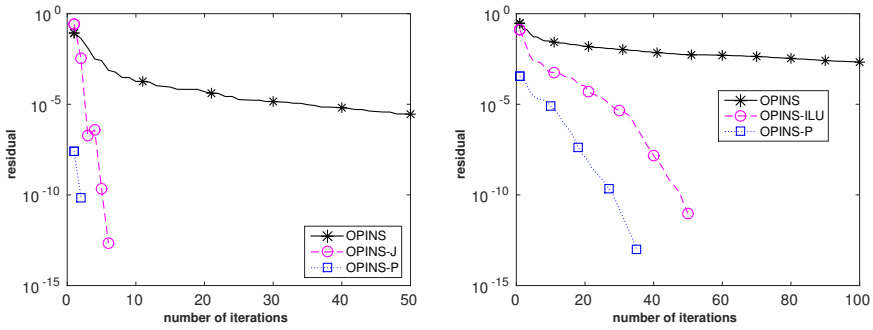

(a) 3d-var (b) sherman5

Figure 2: Convergence history of relative residual in $\boldsymbol{x}$ of OPINS versus preconditioned OPINS for 3d-var and sherman5.

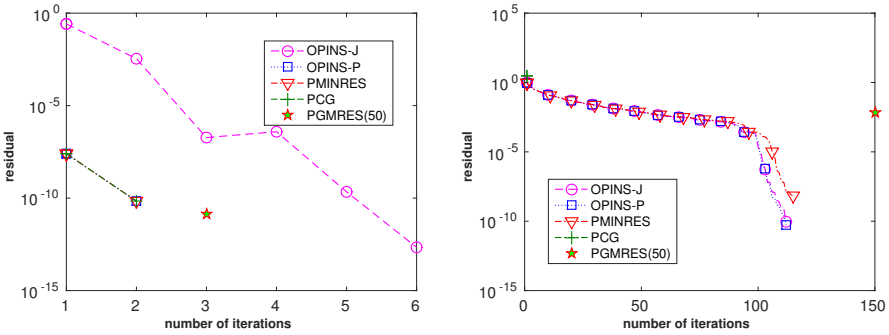

(a) 3d-var. (b) random.

Figure 3: Convergence history of relative residuals in $\boldsymbol{x}$ for OPINS and projected Krylov methods for 3d-var and random.

systems, we only consider OPINS-J and OPINS-P. We compare them against other Krylov subspace methods, including PMINRES and PCG with constraint preconditioners, which are equivalent to implicit null-space methods and are effective choices for the saddle-point systems that we are considering. As a reference, we also consider the preconditioned GMRES with the constraint preconditioner applied to the original system (1). The number of restarts is set to 50. We calculated the error for GMRES only after the solver has converged, so only one data point is available for it.

We solve three test problems, mosarqp1, 3d-var and random. Figure 3(a) and (b) show the convergence of the residual in terms of $\boldsymbol{x}$ for 3d-var and random. In 3d-var, all methods perform well with OPINS-P being similar to PMINRES. In terms of solution accuracy, all strategies give accurate $[\boldsymbol{x}, \boldsymbol{y}]$ as shown in Table 2.

In the random test, the $\boldsymbol{A}$ block is a 100×100 symmetric, nonsingular and indefinite matrix. $\boldsymbol{B}$ is a randomly generated 20×100 dense matrix with full

Table 2: Relative residual in $[\boldsymbol{x}, \boldsymbol{y}]$ upon convergence for OPINS and convergence or stagnation of projected Krylov methods.

Problem	OPINS-J	OPINS-P	PMINRES	PCG	PGMRES(50)
3d-var	7.8e-13	2.5e-10	2.5e-10	2.5e-10	1.3e-11
mosarqp1	2.1e-11	3.9e-11	1.3e-10	1.3e-10	5.4e-11
random	1.2e-12	1.2e-12	3.5e-4	break	[1]1.5e-4

rank. Figure 3(b) shows the $\boldsymbol{x}$ errors of different strategies. In this example, $\boldsymbol{A}$ is no longer positive definite, and PCG breaks at the initial iterations. Compared to PGMRES, OPINS and PMINRES converged much faster both in terms of $\boldsymbol{x}$ and in terms of $[\boldsymbol{x}, \boldsymbol{y}]$. One reason is that OPINS and PMINRES are symmetric methods that can take advantage of a complete Krylov subspace basis. On the other hand, the $\boldsymbol{x}$ residual of PMINRES will stagnate around 10^{-8}. This is due to the instability of the algorithm as mentioned in [19]; see Subsection 5.2.1 for a more detailed analysis and comparison. Due to its stability and faster convergence, OPINS delivered the most accurate solution among all the methods both in terms of $\boldsymbol{x}$ and in terms of $[\boldsymbol{x}, \boldsymbol{y}]$.

5.2.1 Stability of OPINS versus PMINRES

In this section, we study the stability of OPINS and PMINRES. The difference between PCG and PMINRES with a constraint preconditioner mainly lies in the underlying solver. Since MINRES is more robust than CG, we restrict our attention to PMINRES. In our previous discussion, PMINRES and OPINS with the projected preconditioner are both equivalent to a preconditioned null-space method. However PMINRES may stagnate for some problems, as shown in Figure 3 (b).

To further illustrate this, we consider the random test used in Subsection 5.2. OPINS with the two preconditioners are applied here. PMINRES uses the constraint preconditioner with $\boldsymbol{G}$ chosen as the diagonal part of $\boldsymbol{A}$. To examine the stability of PMINRES, we consider PMINRES+IR(0) and PMINRES+IR(1), which denote no iterative refinement and one step of iterative refinement, respectively. The constraint preconditioner is solved by factorization. Figure 4(a) shows the convergence of various strategies. It can be observed that OPINS-J, OPINS-P and PMINRES+IR(1) have similar convergence behaviors. On the other hand, PMINRES failed to converge to the desired tolerance if no iterative refinement is applied. Another example is the fracture problem. This system is singular, but the projected MINRES is also applicable. We set $\boldsymbol{G}$ to be the identity matrix in the constraint preconditioner. No preconditioner is used in OPINS. From Figure 4(b), we can see that OPINS is similar to the more stable PMINRES+IR(1), while the residual of PMINRES+IR(0) oscillates.

In addition, OPINS is more stable compared to PMINRES with iterative refinement, because the latter may still stagnate. We consider the example

[1]This is the error for PGMRES after 150 iterations.

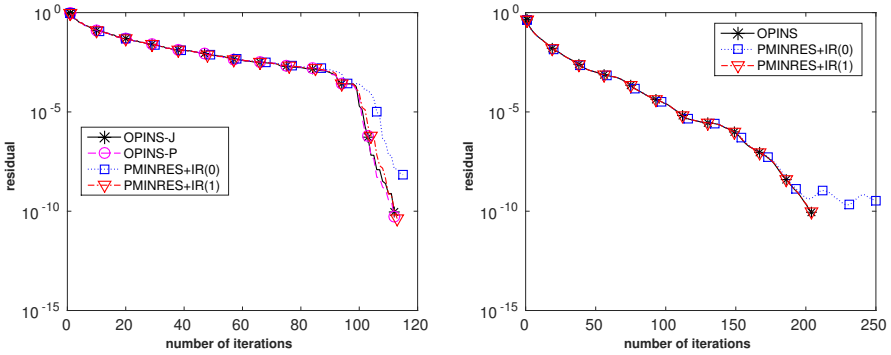

(a) Relative residuals in $\boldsymbol{x}$ for random. (b) Relative residuals in $\boldsymbol{x}$ for fracture.

Figure 4: Comparison of OPINS and projected MINRES with/without iterative refinement for random and fracture.

can_61 from the matrix market database [8]. The matrix is set as the $\boldsymbol{A}$ part of the saddle-point system. $\boldsymbol{B}$ is a randomly generated 20×61 dense matrix with full rank. In this example, we include an additional strategy, PMINRES+IR(2) which uses two steps of iterative refinement per iteration. Figure 5(a) displays the residuals of the five strategies. Both OPINS-J and OPINS-P converged to the specified accuracy. However, PMINRES+IR(0) and PMINRES+IR(1) both stagnated early. Even additional iterative refinement could not increase the accuracy further as PMINRES+IR(2) still stagnated.

One reason for the stagnation of PMINRES is that $\boldsymbol{x}_n$ in the constraint preconditioner may deviate from null $(\boldsymbol{B})$ during the iteration. Eventually the large residual in $\boldsymbol{B}\boldsymbol{x}_n$ may cause MINRES to stagnate [19]. This is evident in Figure 5(b), where we plot $\|\boldsymbol{B}\boldsymbol{x}_n\|$ and $\left\|\boldsymbol{B}\boldsymbol{\Pi}_{\boldsymbol{U}}^{\perp}\boldsymbol{w}\right\|$ for PMINRES and OPINS. For PMINRES with and without iterative refinement, $\|\boldsymbol{B}\boldsymbol{x}_n\|$ grew, and the growth of non null-space component was significant even with two steps of iterative refinement. In contrast, this norm stayed at close to machine precision for OPINS throughout the computation, because OPINS explicitly projects the solution onto null $(\boldsymbol{B})$. For this reason, OPINS is more stable than PMINRES.

5.3 Assessment for Symmetric Singular Systems

When solving singular systems, no preconditioners are used to avoid altering the minimum-norm solution. The methods for comparison are CG, MINRES and GMRES with the same constraint preconditioner. For GMRES, we calculate the error after the solver has converged.

In this subsection, we consider some test cases with singular systems. The test problem fracture, comes from a finite element analysis of nonlinear elasticity model with cracks. In this system, $\boldsymbol{A}$ represents the global stiffness matrix. The constraint matrix $\boldsymbol{B}$ is used to enforce various boundary conditions, such as

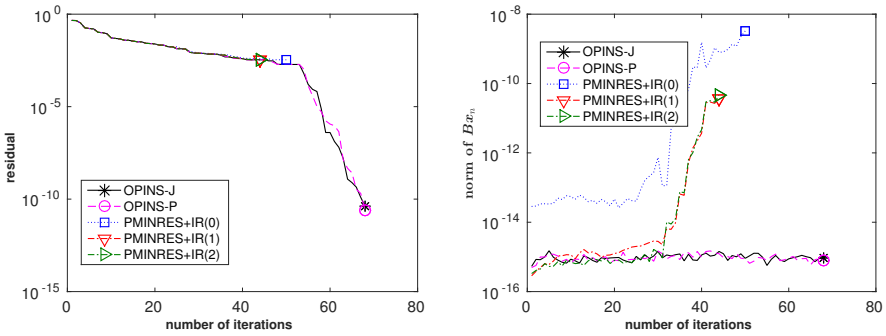

(a) Convergence history of relative residual in $\boldsymbol{x}$.

(b) History of $\|\boldsymbol{B}\boldsymbol{x}_n\|$ in OPINS and PMINRES.

Figure 5: Comparison of OPINS and projected MINRES with/without iterative refinement for `can_61`, for which iterative refinement could not improve accuracy for projected MINRES due to loss of orthogonality.

sliding boundary conditions and contact constraints. Due to the cracks opening along edges, some pieces are isolated completely from the main body. This introduces additional singularity to $\boldsymbol{A}$ such that $\text{null}(\boldsymbol{A}) \cap \text{null}(\boldsymbol{B}) \neq \{0\}$. In addition, $\boldsymbol{B}$ contains redundant constraints, and it is rank-deficient. We apply OPINS without preconditioners to solve the problem. For PMINRES, PCG and PGMRES, $\boldsymbol{G}$ is set to be the identity in the constraint preconditioner. Since $\boldsymbol{B}$ is rank-deficient, an extra step of QR factorization is used to remove linearly dependent rows.

5.3.1 Convergence Comparison

Figure 6 shows the convergence of strategies. It can be seen that OPINS converges monotonically to the desired accuracy while the errors of PMINRES and PCG oscillate. In particular, the error of PCG starts to grow after a certain number of iterations. This shows that CG is not as robust as MINRES for solving singular systems. In terms of the overall residual, OPINS has around 10^{-9} while the residuals of PCG and PMINRES stay around 1. The reason why PCG and PMINRES have large overall errors is that they do not necessarily minimize the error of $[\boldsymbol{x}, \boldsymbol{y}]$. To get a more accurate $\boldsymbol{y}$, we need to solve the least-squares problem $\boldsymbol{B}^T\boldsymbol{y} = \boldsymbol{f} - \boldsymbol{A}\boldsymbol{x}$ once convergence in $\boldsymbol{x}$ is reached. On the other hand, the residual of GMRES is about 10^{-6}. However the converged solution is incorrect, since $\|\boldsymbol{g} - \boldsymbol{B}\boldsymbol{x}\|$ is around 1. Overall OPINS is more stable and accurate than the other approaches.

5.3.2 Solution Norm

To illustrate the minimum norm property of OPINS, we compare OPINS with truncated SVD. The $\boldsymbol{x}$ and $\boldsymbol{y}$ components are measured by 2-norm. The relative

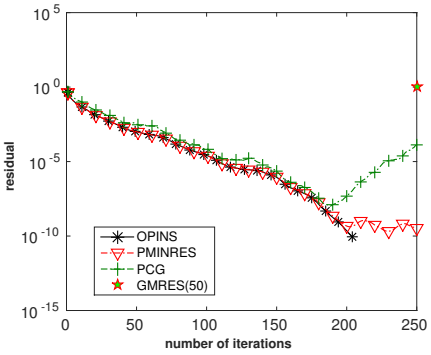

Figure 6: Convergence history of relative residual in x for OPINS and Krylov methods for singular system from fracture.

2-norm of the residuals are provided as references. In the fracture system, large Lamè parameters are used in calculating element stresses. As a result, there is a 10^{10} gap between the scaling of $\boldsymbol{A}$ and $\boldsymbol{B}$ which makes the system very ill-conditioned. The conditioning will be improved if we scale $\boldsymbol{A}$ and $\boldsymbol{f}$ by 10^{-10}. After the operation, $\boldsymbol{x}$ should remain unchanged while $\boldsymbol{y}$ will be scaled by 10^{-10} accordingly. We apply OPINS and truncated SVD (TSVD) to both scaled and unscaled systems, where the tolerance for TSVD was 10^{-12}. As shown in Table 3, OPINS is stable with and without scaling. Its $\boldsymbol{x}$ component stays unchanged, and it has the same norm as the TSVD solution on the scaled system. Without scaling, TSVD produced a wrong solution that is far away from the constraint hyperplane. This is evident from Table 4, where $\|\boldsymbol{g}-\boldsymbol{B}\boldsymbol{x}\|$ is very large for TSVD. In contrast, OPINS finds the minimum-norm $\boldsymbol{x}$ independently of the scaling. Therefore, it is advantageous to solve $\boldsymbol{x}$ and $\boldsymbol{y}$ separately over solving them together, and it is not advisable to use any linear solver, including TSVD or GMRES, as black box solvers for saddle-point problems.

In terms of the $\boldsymbol{y}$ component, the solution of OPINS has a slightly larger norm than the TSVD solution for the scaled system. This is due to the rank deficiency of $\boldsymbol{B}$. For the case where $\boldsymbol{B}$ has full rank, we consider a random system. $\boldsymbol{A}$ is a 100×100 dense and semi-definite matrix with rank $= 50$. $\boldsymbol{B}$ is a randomly generated 20×100 dense matrix with full rank. By construction, the system is singular with rank equal to 90. Since $\boldsymbol{A}$ is semi-definite and $\boldsymbol{B}$ has full rank, we can expect both $\boldsymbol{x}$ and $\boldsymbol{y}$ components of the OPINS solution to have minimum norms, which is evident in Table 3.

6 Conclusions

In this paper, we introduced a new implicit null-space method, called OPINS, for solving saddle-point systems, especially for applications with relatively few constraints compared to the number of solution variables. These systems may

Table 3: Norms of solutions of OPINS versus truncated SVD.

Test Problems	OPINS		TSVD	
	$\|\boldsymbol{x}\|$	$\|\boldsymbol{y}\|$	$\|\boldsymbol{x}\|$	$\|\boldsymbol{y}\|$
fracture	1.53×10^{-4}	1.79×10^{7}	2.86×10^{-5}	4.25×10^{-11}
scaled fracture	1.53×10^{-4}	1.79×10^{-3}	1.53×10^{-4}	1.54×10^{-3}
random-s	5.56	2.41	5.56	2.41

Table 4: Errors of solutions of OPINS versus truncated SVD.

Test Problems	OPINS		TSVD	
	$\|\boldsymbol{g} - \boldsymbol{B}\boldsymbol{x}\|$	$\|\boldsymbol{r}\|$	$\|\boldsymbol{g} - \boldsymbol{B}\boldsymbol{x}\|$	$\|\boldsymbol{r}\|$
fracture	1.0×10^{-17}	1.2×10^{-9}	1.1	1.9×10^{-10}
scaled fracture	8.6×10^{-18}	5.6×10^{-10}	1.1×10^{-14}	1.3×10^{-13}
random-s	3.7×10^{-16}	2.1×10^{-13}	2.3×10^{-14}	4.6×10^{-15}

be nonsingular, or be singular but compatible in terms of $\boldsymbol{f}$. Instead of finding the null space of the constraint matrix explicitly, OPINS uses an orthonormal basis of its orthogonal complementary subspace to reduce the saddle-point system to a singular but compatible system. We showed that OPINS is equivalent to a null-space method with an orthonormal basis for nonsingular systems. In addition, it can solve singular systems and produce the minimum-norm solution, which is desirable for many applications. Because of its use of orthogonal projections, OPINS is more stable than other implicit null-space methods, such as the projected Krylov methods. Despite its core equation is singular, its nonzero eigenvalues have the same distribution as that of the null-space method with an orthonormal basis, so an iterative solver would converge at the same rate. We proposed effective preconditioners for OPINS, based on the work for projected Krylov methods. For singular saddle-point problems, one should only use left preconditioners that do not alter the null space of the coefficient matrix, to ensure the minimum-norm solution. A future research direction is to develop more effective preconditioners for singular saddle-point systems. Another direction is the parallelization of OPINS, which will require a parallel rank-revealing QR factorization, such as that in [12], as well as parallel KSP solvers, such as those available in PETSc [2].

References

[1] M. F. Adams, *Algebraic multigrid methods for constrained linear systems with applications to contact problems in solid mechanics*, Numer. Linear Algebra Appl., 11 (2004), pp. 141–153.

[2] S. Balay, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang, *PETSc*

users manual, Tech. Rep. ANL-95/11 - Revision 3.0.0, Argonne National Laboratory, 2008.

[3] R. E. Bank, B. D. Welfert, and H. Yserentant, *A class of iterative methods for solving saddle point problems*, Numer. Math., 56 (1989), pp. 645–666.

[4] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. V. der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, 2nd Edition*, SIAM, Philadelphia, PA, 1994.

[5] K. J. Bathe, E. Ramm, and E. L. Wilson, *Finite element formulations for large deformation dynamic analysis*, Int. J. Numer. Meth. Engrg., 9 (1975), pp. 353–386.

[6] M. Benzi, G. H. Golub, and J. Liesen, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137.

[7] M. Benzi and A. Wathen, *Some preconditioning techniques for saddle point problems*, in Model Order Reduction: Theory, Research Aspects and Applications, W. H. A. Schilders, H. A. van der Vorst, and J. Rommes, eds., vol. 13 of Mathematics in Industry, Springer, 2008, pp. 195–211.

[8] R. F. Boisvert, R. Pozo, K. Remington, R. F. Barrett, and J. J. Dongarra, *Matrix market: A web resource for test matrix collections*, in Proceedings of the IFIP TC2/WG2.5 Working Conference on Quality of Numerical Software: Assessment and Enhancement, London, UK, 1997, Chapman & Hall, Ltd., pp. 125–137.

[9] J. H. Bramble and J. E. Pasciak, *A preconditioning technique for indefinite systems resulting from mixed approximations of elliptic problems*, Math. Comput., 50 (1988), pp. 1–17.

[10] T. F. Chan, *Rank revealing QR factorizations*, Linear Algebra Appl., 88 (1987), pp. 67–82.

[11] S.-C. T. Choi, C. C. Paige, and M. A. Saunders, *MINRES-QLP: A Krylov subspace method for indefinite or singular symmetric systems*, SIAM J. Sci. Comput., 33 (2011), pp. 1810–1836.

[12] T. A. Davis, *Algorithm 915, SuiteSparseQR: Multifrontal multithreaded rank-revealing sparse QR factorization*, ACM Trans. Math. Softw., 38 (2011), p. 8.

[13] E. de Sturler and J. Liesen, *Block-diagonal and constraint preconditioners for nonsymmetric indefinite linear systems. Part I: Theory*, SIAM J. Sci. Comput., 26 (2005), pp. 1598–1619.

[14] D. Fong and M. M. Saunders, *LSMR: An iterative algorithm for sparse least-squares problems*, SIAM J. Sci. Comput., 33 (2011), pp. 2950–2971.

[15] T.-P. Fries, A. Byfut, A. Alizada, K. W. Cheng, and A. Schröder, *Hanging nodes and XFEM*, Int. J. Numer. Meth. Engrg., 86 (2011), pp. 404–430.

[16] J. R. Gilbert and M. T. Heath, *Computing a sparse basis for the null space*, SIAM J. Alg. Disc. Meth., 8 (1987), pp. 446–459.

[17] G. H. Golub and C. F. Van Loan, *Matrix Computations*, Johns Hopkins, 4th ed., 2013.

[18] N. I. M. Gould, M. E. Hribar, and J. Nocedal, *On the solution of equality constrained quadratic programming problems arising in optimization*, SIAM J. Sci. Comput., 23 (2001), pp. 1376–1395.

[19] N. I. M. Gould, D. Orban, and T. Rees, *Projected krylov methods for saddle-point systems*, SIAM J. Matrix Anal. & Appl., 35 (2014), pp. 1329–1343.

[20] A. Greenbaum and L. Gurvits, *Max-min properties of matrix factor norms*, SIAM J. Sci. Comput., 15 (1994), pp. 348–358.

[21] J. Huang, *Constrained Variational Analysis Integrating Vertical and Temporal Correlations*, PhD thesis, Applied Math and Statistics, Stony Brook University, 2012.

[22] C. Keller, N. I. M. Gould, and A. J. Wathen, *Constraint preconditioning for indefinite linear systems*, SIAM J. Matrix Anal. & Appl., 21 (2000), pp. 1300–1317.

[23] I. Maros and C. Mészáros, *A repository of convex quadratic programming problems*, Optim. Method. Softw., 11 (1999), pp. 671–681.

[24] J. Nocedal and S. Wright, *Numerical Optimization*, Springer Science & Business Media, 2006.

[25] C. C. Paige and M. A. Saunders, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.

[26] C. C. Paige and M. A. Saunders, *LSQR: An algorithm for sparse linear equations and sparse least squares*, ACM Trans. Math. Softw., 8 (1982), pp. 43–71.

[27] L. Reichel and Q. Ye, *Breakdown-free GMRES for singular systems*, SIAM J. Matrix Anal. Appl., 26 (2005), pp. 1001–1021.

[28] Y. Saad, *Iterative methods for sparse linear systems*, SIAM, 2nd ed., 2003.

[29] Y. Saad and M. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869.

[30] J. Schöberl and W. Zulehner, *Symmetric indefinite preconditioners for saddle point problems with applications to PDE-constrained optimization problems*, SIAM J. Matrix Anal. & Appl., 29 (2007), pp. 752–773.

[31] M. Tuma, *A note on the LDL^T decomposition of matrices from saddle-point problems*, SIAM J. Matrix Anal. & Appl., 23 (2002), pp. 903–915.

[32] S. A. Vavasis, *Stable numerical algorithms for equilibrium systems*, SIAM J. Matrix Anal. & Appl., 15 (1994), pp. 1108–1131.